LOS ALAMOS NATIONAL LABORATORY 1997 R&D 100 AWARD WINNER
LOS ALAMOS NATIONAL LABORATORY 1997 R&D 100 AWARD WINNER

RN97-093-007

The High Performance Storage System is designed to make data immediately available to all networked computing nodes, regardless of size, distribution, or location. It is highly scalable in capacity and performance so that the storage system can grow incrementally as needs increase.

High Performance Storage System

Los Alamos National Laboratory, Lawrence Livermore National Laboratory, Oak Ridge National Laboratory, Sandia National Laboratories, and IBM Global Government Industry

Today's computer models and simulations, digitized information libraries, and high-energy physics data require an immense amount of electronic storage space. The High Performance Storage System (HPSS)—developed by Los Alamos National Laboratory, Lawrence Livermore National Laboratory, Sandia National Laboratories, Oak Ridge National Laboratory, and IBM Global Government Industry—is a new, hierarchical, software data-storage system designed to handle the storage capacity and throughput requirements of the most demanding supercomputers and high-speed networks.

The Invention—Characteristics and Advantages

The most innovative characteristic of HPSS is the scalable architecture that supports very-high-performance input/output and very large stores and libraries. The key architectural breakthroughs are its network-centered design, parallel input/output, and servers that can be distributed and replicated.

HPSS is faster and has more storage capacity than any other storage system, allowing users to manage petabytes—1,000 trillion bytes—of data. Simply storing huge amounts of data is not all that's required for data management; users must be able to retrieve data quickly as well. HPSS allows users to transfer data at rates of up to one gigabyte—one billion bytes—per second. Data transfer rates and storage capacities can be scaled to fit users' needs and to avoid bottlenecks. In this way, HPSS allows users unprecedented range and flexibility.

Unlike traditional data-storage systems, HPSS software components are distributed on a network and are not confined to a centralized storage computer. The system implements parallel input/output for disks and tapes, so the amount of data going through the system isn't limited to the fastest single device in the storage system but scales linearly with the number of such devices.

The basic infrastructure of HPSS is the Distributed Computing Environment because of its wide adoption and nearly universal acceptance by the computer industry. The network itself can use any technology supported by standard protocols such as TCP/IP or IPI-3. Users can access HPSS files through any industry-standard FTP or NFS interface, which means any computer platforms that support these protocols can access HPSS.

HPSS accommodates the simultaneous transfer of parallel data streams from multiple storage devices to computers exercising parallel applications—the resulting throughput is therefore the aggregate of the total number of storage devices. Using today's technology, total throughput rates into the gigabyte-per-second range are possible. As the speeds and capacities of storage media and devices increase, HPSS will easily assimilate them to produce transfer rates in the range of multiple gigabytes per second as will be required by parallel computing, imaging, and visualization techniques.

HPSS functions are implemented by various server processes. Most servers can have several copies of a file operating simultaneously—on as many network nodes as necessary—to provide the level of performance and scalability required by an installation. The configuration of the system, including the number and type of operating servers, is controllable in real-time by the operator. This represents a major paradigm shift away from a single storage computer to a scalable array of network-attached storage nodes, devices, and servers.

HPSS uses advanced techniques such as transaction management to provide security and protection of vast amounts of data moving at high speed. "Transactions" are related groups of functions that must take place together to maintain data integrity and reliability. HPSS makes sure that all parts of a transaction are

LOS ALAMOS NATIONAL LABORATORY 1997 R&D 100 AWARD WINNER
LOS ALAMOS NATIONAL LABORATORY 1997 R&D 100 AWARD WINNER
LOS ALAMOS NATIONAL LABORATORY 1997 R&D 100 AWARD WINNER

completed, or the entire transaction is "backed out" to the original state if a failure occurs. The use of the advanced techniques also ensures that data is protected from unauthorized use and from corruption. HPSS ensures that only authenticated users and nodes gain access to data permitted to them.

Applications

HPSS was designed to provide very-high-speed transfer of the massive quantities of data generated by supercomputers, massively parallel computers, or clusters of workstations and data-collection and imaging devices. The national laboratories developing HPSS are dependent upon the capabilities that HPSS provides. In particular, the striking storage and access requirements of the Accelerated Strategic Computing Initiative (ASCI) project, focused at these laboratories, can only be met with HPSS. For example, the ASCI project is implementing an environment that will scale to over 100 teraflops (100 trillion floating-point operations per second) with projected storage capacities in the exabyte (one million terabytes) range.

HPSS is valuable in any computing system or installation emphasizing large quantities of data and requiring very high transfer rates. It is crucial for large modeling and simulation applications with large storage requirements. HPSS can be used to archive and move huge amounts of information for telemedicine applications; banking, finance and insurance; fingerprinting systems; the Human Genome Project; animation and entertainment; high-energy physics, astronomy, and other scientific applications; oil exploration; and other areas that require storage and movement of data.

In September 1996, IBM announced the commercial availability of HPSS. In addition to the HPSS development partners, many other organizations that use high-performance computers and therefore have large, high-end data-storage needs are implementing HPSS. These organizations include the Cornell Theory Center, the San Diego Supercomputer Center, the Maui High Performance Computing Center, NASA Langley Research Center, the Fermi National Accelerator Laboratory, the University of Washington, and the Jet Propulsion Laboratory (California Institute of Technology).

Recognizing the need to develop vastly improved storage systems and recognizing the fact that no single organization could address all the problems, staff members in the Data Storage Systems Group at Los Alamos joined in 1993 with computer scientists from the four other collaborating organizations to develop HPSS.

Los Alamos has been actively involved in developing data-storage systems for over 20 years. This development activity has produced several data storage systems, including the High Performance Data System, which was an immediate precursor to HPSS. As a result of such development efforts, Los Alamos was able to bring both experience and software to the development of HPSS. Los Alamos actively participates in the design and development of HPSS, as well as in the technical and executive oversight of the project.

Left to right: Ling-Ling Chen, Tyce T. McLarty, D. Lynn Kluegel, Cheryl A. Ramsey, John W. Blaylock, Lynn K. Jones, Danny P. Cook, Bart Parliman. Not pictured: Ronald D. Christman, Bill Collins, Gary A. Grider, Frank A. Maestas, Mark A. Roschke.

For more information about HPSS, please contact the Civilian and Industrial Technology Program Office, Los Alamos National Laboratory, P.O. Box 1663, Mail Stop C331, Los Alamos, NM 87545. Telephone (505) 665-9090, Fax (505) 667-0603.

LALP-97-146 August 1997

Los Alamos
NATIONAL LABORATORY
A US Department of Energy Laboratory